Service Oriented Traffic Models for Future Backbone Networks

Eleni Palkopolou, Dominic A. Schupke, Claus G. Gruber, Andreas Kirstädter
Siemens Networks GmbH & Co KG, 81730 Munich, Germany
Christian Merkle
Technische Universität München, 80333 Munich, Germany, christian.merkle@tum.de

Abstract

In this paper we present a novel approach to assess the impact of new and existing services on traffic volume in current and future backbone networks. Several proposals to model traffic load in access- and backbone networks exist in the literature. These proposals consider current Internet traffic like http, smtp, ftp, and Peer to Peer (P2P). We expect, however, that there will be a change in traffic load for future networks caused by services like IP Television (IPTV), Video on Demand (VoD), and Virtual Private Networks (VPN). Additionally, population-based models may no longer be applicable due to the widespread of service-providers and hierarchical routing through network peering points. Therefore, it is important to reassess future traffic volumes and traffic patterns and to identify those services that have the most impact on the networks. We model today's traffic volume of each of the described services and estimate future traffic volumes taking peering points into account. To illustrate the different traffic flows and to characterize the traffic distribution we apply our results to a Germany reference network.

1 Introduction

The last years have already seen an enormous growth of bandwidth needs in backbone networks of around 40% per year. Recently, this growth has been sharply accelerated by the market acceptance of new business services and it has been fuelled by the roll-out of more and more high-speed x Digital Subscriber Line (xDSL)-based residential access technologies. The latter together with flat-rate business models are enabling more and more traffic of Peer-to-Peer (P2P) applications and all the emerging community-oriented end-user applications (video file sharing, blogging, gaming, second life etc.) in the Internet that are commonly summarized under the keyword Web2.0. Thus, traffic growth rates up to 100% per year are anticipated for the backbone networks - putting an end to all rumors about a bandwidth glut in the backbone networks.

At the same time, the revenues of the carriers stay pretty constant or are only rising in the order of a few percent per year. It is commonly understood that the resulting steep decline in revenue per bandwidth unit can only be absorbed via the introduction of new network architectures and technologies together with suitable network planning.

Thus the chance to substantially improve the cost situation of operators is also depending on the availability of traffic models for backbone networks that take into account the new service developments and that are able to predict their future bandwidth requirements.

Another important aspect is the basic traffic flow patterns in the current and future Internet: Peering and transport relations as well as content delivery architectures strongly bias load models into asymmetric configurations around certain hubs in the backbone networks. A classical example here is P2P traffic - already dominating Internet usage to a large extent. While the endpoints of P2P relations might well be dispersed across countries and networks, the traffic flows themselves are always confined to pass through the Internet Exchange Points (IXPs) interconnecting the networks of the different Internet Service Providers (ISPs).

To account for these service-oriented patterns and developments, this paper takes a novel approach for modeling backbone network traffic. We first (in Section 3) consider the different services and evaluate the typical traffic load patterns they impose on the backbone in order to be able to rank their importance. In the next step we then consider the traffic matrices generated by each service - putting strong emphasis on the asymmetry generated by routing, peering, and hosting. Finally we then parametrize and then aggregate the influences of the single services to obtain a scalable traffic model for future backbone networks. Section 4 shows some results on this.

2 Related Work

Reference [1] analyzes P2P traffic characteristics and shows that although the majority of the shared files have the size of typical song files, most of the traffic is generated by movie downloads. The results in this paper conclude that 20% of the files account for more then 80% of the downloads. In [2] a traffic model for a US optical network is presented. The traffic growth rates from voice traffic, transaction data traffic, and Internet traffic along with the requirements for the network are analyzed. With these growth rates a traffic model is developed and the link capacities needed

for the year 2004 are calculated. In [3] a monitoring system that is able to measure packet-level streams on backbone links is described. It is shown that link load characteristics vary from link to link and are often correlated to the nature of the customers connected to the point of presence. It is also shown that some links no longer have web traffic as dominant traffic. In [4] the traffic parameters based on a population model for three different networks are derived. The paper presents static and dynamic traffic characteristics for the different optical networks.
These approaches elaborate on the traffic volume and traffic characteristics of a set of known services. However, new services, like IPTV, VoD, and user generated content, with new traffic characteristics and higher data volume become more and more important and lead to higher demands in the backbone network.

3 Services

In this section the characteristics of IPTV, VoD, P2P, user generated content, and VPNs are presented. The traffic flow of each service is illustrated in a figure, which depicts a 17 node Germany reference network from [4]. The arrows indicate the direction of each traffic flow and the number of arrows is related to the amount of traffic, which is routed over the link.

3.1 IP Television

The telecommunications sector is rapidly evolving to offer commercially live broadcast TV over IP - known as Internet Protocol TV (IPTV). Scaling to mass markets necessitates a reliable and cost efficient network infrastructure all the way from the central head ends where the video is sourced to the customers [5]. Thus, in order to achieve the desired efficiency, "broadcast" IPTV is delivered over a multicast distribution scheme (Multicast Distribution Tree) in the backbone, avoiding multiple transmissions of the same content over the network.
Throughout this paper, the source of all IPTV content is denoted as the Super Hub Office (SHO) [5]. For redundancy reasons two SHOs are needed to ensure reliable transmission in case of catastrophic failure of one of the SHOs. Video streams, transmitted from the SHO, are received at the Video Hub Offices (VHOs) and in turn transmitted to the customers. Consequently, the traffic on the backbone is independent of the actual number of IPTV subscribers. At the VHOs, content can be stored locally or even processed before it is transmitted to the users.
In the following, the resulting traffic patterns will be studied, using as a reference network the 17 node Germany network [4]. The SHOs are located in Frankfurt, where the dominant Internet Exchange Point (DE-CIX) exists, and in Munich. However, at each time instant only one SHO is active and the remaining nodes act as the VHOs, as depicted in **Figure 1**. The traffic on every link that is part of the multicast tree is determined by the number of IPTV providers and the offered channels as well as on the codecs and the resolution used. Using the MPEG4 codec the required data rate for Standard Definition TV (SDTV) is in the range of $3.5\ \mathrm{Mbit/s}$ to $5\ \mathrm{Mbit/s}$ and for High Definition TV (HDTV) it is $8\ \mathrm{Mbit/s}$ to $12\ \mathrm{Mbit/s}$.

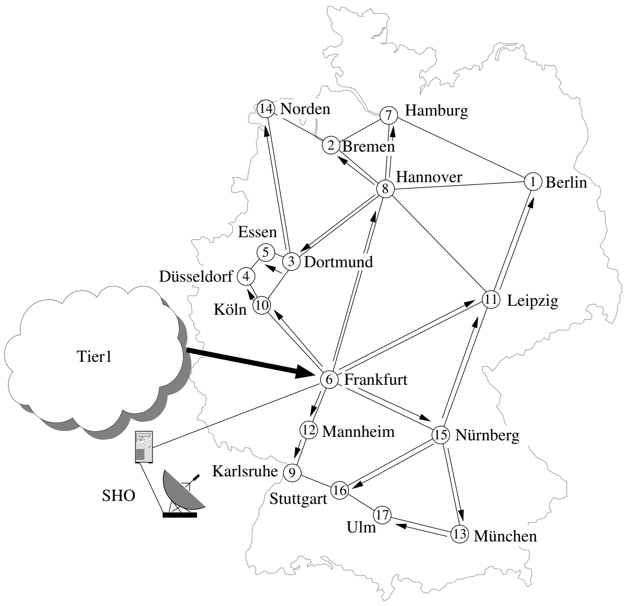

Figure 1: IP Television traffic pattern

By the term IPTV, many related supporting service offerings may be implied. Considering the broadcast service, the offerings of broadcast TV and audio channels as well as Near Video on Demand (nVoD/Pay-Per-View) can be included. The term nVoD refers to the service that enables the user to choose a program from a predefined selection that is broadcast at fixed time intervals. The on-demand service includes the classical offerings of Video on Demand (VoD), Music on Demand, Network Personal Video Recorder (nPVR) and Time-Shift TV. In addition to the broadcast and the on-demand service offerings, interactive services may be bundled in the IPTV package, including among others interactive information, interactive TV, and online gambling. In the scope of this paper, only the bandwidth driving applications are examined in more detail. The term IPTV refers exclusively to the broadcast service, while the on-demand service is examined separately in the Video on Demand section.
The traffic volume is estimated for the 17 node Germany network presented in Figure 1. After researching the current market status, we assume two IPTV service providers, each offering 100 channels, consisting of a mixture of SDTV and HDTV channels. Considering that there will be a shift toward HDTV that promises a higher quality viewing experience, in order for IPTV to remain competitive in the service offerings, the worst case in terms of bandwidth consumption is examined.
This results in an aggregated traffic volume of $2.4\ \mathrm{Gbit/s}$ per link used, corresponding to a total traf-

fic value of 38.4 Gbit/s with the routing displayed in Figure 1. Projecting the future traffic volume (2010), and assuming 400 HDTV channels each at 12 Mbit/s, results in an aggregated traffic volume of 4.8 Gbit/s per link used. This corresponds to a total traffic value of 78.6 Gbit/s with the routing displayed.

3.2 Video on Demand

VoD is an additional service, usually combined with the offering of IPTV broadcast channels. As the name suggests, it is the delivery of video content to the user upon request. VoD content is sent to each individual user as a real-time dedicated stream. In order to minimize the amount of traffic that must be carried across the backbone, the popular VoD content is stored locally at the VHOs.

When new content is to be sent to the VHOs we assume, that it is pushed from the SHO to the VHOs during off-peak periods [5]. These transfers do not require real-time delivery, and bulk-transfer applications (e.g. ftp) can be used to ensure reliable delivery. Therefore it has minimal impact on the network design and architecture. In case the service provider wishes to offer a vast variety of VoD content, it may not be cost effective to store the entire content at every VHO. Cache management algorithms may be used to increase the hit ratio of the requested VoD content. The hit ratio is defined as the percentage of VoD requests by the users that get served by the local VHO. The remaining requests will be served by the SHO via an unicast delivery scheme across the backbone. As a result the unicast VoD traffic will vary according the achieved hit ratio and the peak usage by the subscribers. At this point two different cases are examined in the next subsections: VoD content stored exclusively in the VHOs and VoD content that is stored both in the SHOs and in the VHOs.

3.2.1 VoD content exclusively stored in the VHOs

In **Figure 2**, the distribution of the VoD content from the SHO to the VHOs is assuming a shortest path algorithm. The actual traffic per span of the shortest path tree depends on the number of the offered on-demand titles, the refresh rate of the titles, and the deployed codecs. In this paper we assume that MPEG4 is used. The refresh rate is the number of videos, which are renewed at local VHOs monthly.

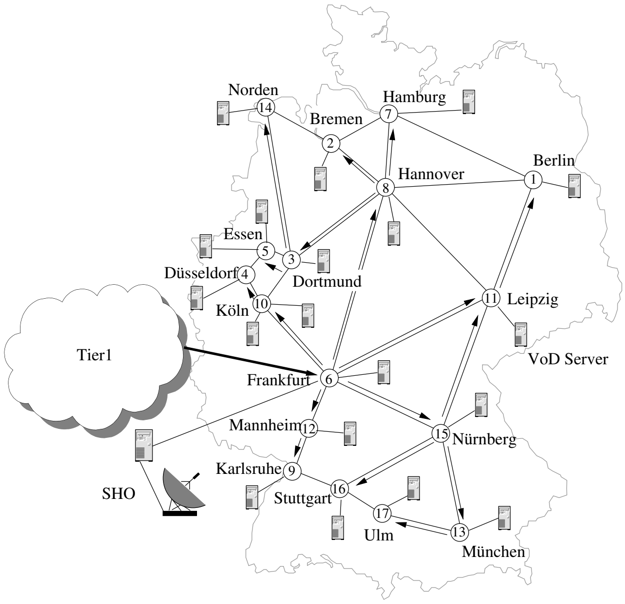

Figure 2: Locally stored Video on Demand traffic pattern

Concidering the content that is locally stored in the VHOs, the traffic volume is estimated for the examined reference network. To calculate the bandwidth utilization per link, we consider 1500 on-demand titles with an average size of 4.5 GB. A worst-case assumption is made for the refresh rate, meaning that they are renewed monthly by 100%. This results in an aggregated traffic volume of 20 Mbit/s per link used. The new content is sent to the VHOs during off-peak periods and it does not require real-time transmission. Thus the impact on the network is minimal as stated previously.

To calculate the future (2010) traffic volume, we assume that there are 5000 on-demand titles available, 12 GB to 20 GB each. A data rate of 10 Mbit/s to 12 Mbit/s is considered. This results in an aggregated traffic volume of 300 Mbit/s per link used, assuming again a refresh rate of 100%. Therefore, the impact of the future traffic caused by VoD is again very low compared to a bandwidths of 40 Gbit/s or 100 Gbit/s for a backbone link in the future.

3.2.2 VoD content stored in the VHOs and SHOs

In **Figure 3**, the distribution of the VoD content from the SHO to the individual customers is being depicted. As analyzed previously, the VoD content is being unicast to the customers over the backbone, resulting in increased traffic demands. In this scenario, the traffic carried over the backbone depends on the number of VoD-subscribers, the achieved hit ratio, and the deployed codecs. The number of VoD-subscribers depends on the deployment of high-speed access infrastructure, a prerequisite to delivering VoD services, and in turn on the market penetration. Currently, to the best of our knowledge, no such offer exists in the German IPTV market.

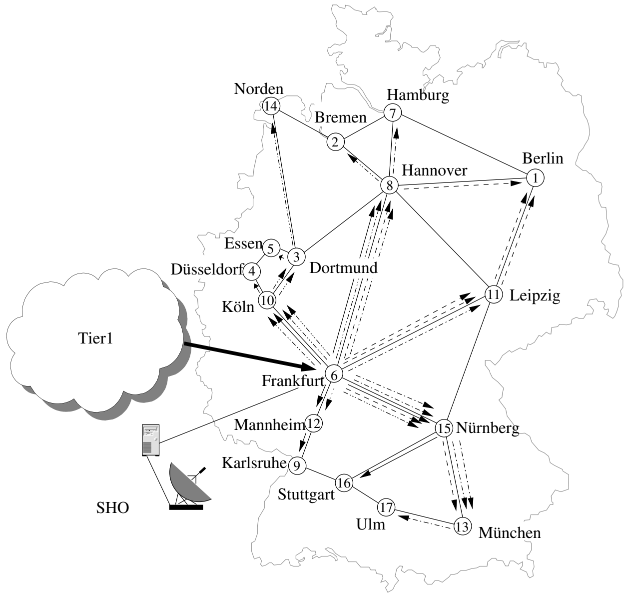

Figure 3: Unicast Video on Demand traffic pattern

Supposing that in 2010 this type of service will be offered, the required data rate on the backbone is calculated. First of all, the percentage of VoD subscribing households is estimated as 10% of the IPTV subscribing households [6], which are projected to reach 2.6 - 2.8 million in 2010 according to Gartner [7]. To calculate the maximum bandwidth needed to deliver this service, we assume that 15% of these subscribers require the service during the busy hour. At this point, the percentage of the peak concurrent subscribers requesting VoD content that is not available in the VHOs must be calculated. The remaining percentage, which corresponds to the subscribers that get served by the local VHOs, is defined as the hit ratio. Hence the traffic imposed on the backbone is directly related to the hit ratio.
In **Figure 4**, the peak total outbound traffic from the SHO is given for different values of the hit ratio.

Hit Ratio in %	Peak Outbound traffic from SHO in Gbit/s
60	201.6
70	151.2
85	75.6
95	25.2

Figure 4: Peak Outbound traffic from SHO

Examining the extreme cases, a hit ratio of 60% results in 201.6 Gbit/s of outbound traffic from the SHO and a hit ratio of 95% in 25.2 Gbit/s of outbound traffic. Thus, the bandwidth utilization to provide these services is much larger than in the scenario before. The selection of the optimal value of the hit ratio is a function of economic and technical reasons, depending on factors such as market penetration and competition.

3.3 Peer to Peer

Over the last years, P2P file sharing systems have evolved to become one of the major traffic contributors in the Internet [8]. P2P networking has emerged as an increasingly popular way for broadband subscribers to share digital content such as music and videos, and an increasingly problematic cause of bandwidth bottlenecks for many broadband providers [9].
Traditional applications like browsing and Email are unidirectional and only active when the user is at the PC. In contrast P2P traffic is symmetrical, since downloads and uploads are concurrently done, and the user does not need to be present during the application's activity time. An issue that comes up, due to the symmetrical nature of P2P traffic, is the caused upstream congestion since networks have typically less capacity provisioned in the upstream direction. Another characteristic of P2P traffic is that it is generally geographically indifferent, since the users can download files from anywhere.
P2P traffic grows with the number of broadband subscribers and their available access speeds. In **Figure 5**, the traffic generated by P2P applications is visualized according to the population of each node.

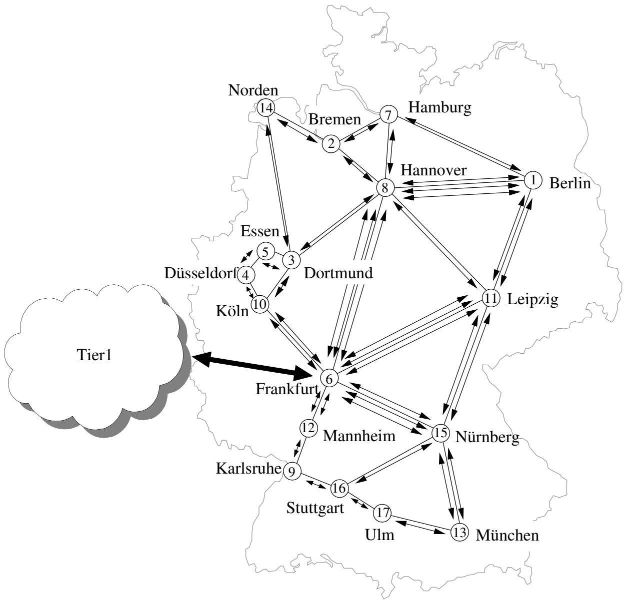

Figure 5: Peer-to-Peer traffic pattern

Almost all of P2P file sharing traffic is international, with percentages higher than 90% in all but a few countries [10]. However, a large portion of German P2P traffic may actually stay within Germany and other German speaking countries.
If we take into account, that Germans show a clear preference to movies either originally in the German language or movies which are later on dubbed into

German, the exchange of video content will be restricted to a large extent in the regions that favor the German language. Since video constitutes 60% to 70% of the P2P traffic [11], the geographical distribution of P2P traffic is expected to show a higher density in German speaking regions.

A study in 2005 showed that 65% of traffic on a service provider's residential broadband network was P2P. This portion of the traffic does not generate extra revenues for the operators and may lead to high peering costs, in case the traffic goes off-net. This constitutes an intense problem in Asia, where some operators have found as much as 80% of broadband traffic leaves the country in search of P2P hosts [9].

In consumer broadband networks 50% to 65% of downstream traffic and 75% to 90% of upstream traffic is P2P [12]. The difference in the percentages of upstream and downstream traffic is due to the asymmetrically provisioned access networks, which generally provide more bandwidth for downloading.

Looking in more detail into the German P2P traffic, a recent study has shown that 30% of daytime and 70% of night-time of the overall Internet traffic in Germany is P2P [11]. As a general remark, most of the international traffic of Germany is routed through DE-CIX, the Frankfurt IXP. Currently there is no sign of decrease of the P2P traffic.

According to various studies [13], [14], there is an untapped locality of up to 86% in P2P workload. This means that 86% of the downloaded content is available within the network, but because the current protocols do not favour neighbouring peers, the content is actually downloaded from external to the network sources. If this locality is exploited, then the potential bandwidth savings would be significant.

Concerning Germany, the absolute data volume has risen by 10% between June and October 2006 [11]. Extrapolating to 2010 leads to a 314% increase in traffic. These values are also in compliance with data from DE-CIX. For the examined reference model, the estimated traffic value for the P2P traffic is $10,539\ \mathrm{Gbit/s}$.

3.4 User Generated Content and Content Delivery Networks

Video search and streaming has seen significant uptake in usage. A survey conducted in late 2005 by the Amsterdam IXP determined that video and audio streaming accounted for 14% of members' Internet traffic, and is expected to be the highest area of growth in future [15].

In order to analyze the traffic pattern of video sharing websites, a representative case study is conducted focusing on the website YouTube. YouTube is the leading net video download site in the US, with 47.7% of the market share of visits to on-line video sites [16]. The site specializes in short, home-made videos. A very important characteristic of its produced traffic is that it is delivered over a unicast scheme. This means that every users request is served independently, resulting in multiple transmissions of the same content over the network. Another characteristic is that the videos are streamed almost instantaneously, imposing strict quality constraints. Currently its servers are located in the US, however its recent acquisition by Google may have an effect on the deployed architecture.

Figure 6 depicts the unicast distribution scheme with the content originating in Frankfurt, since international traffic is mainly routed through Germany's largest IXP. The number of streams per node is proportional to the population and shortest path routing is assumed.

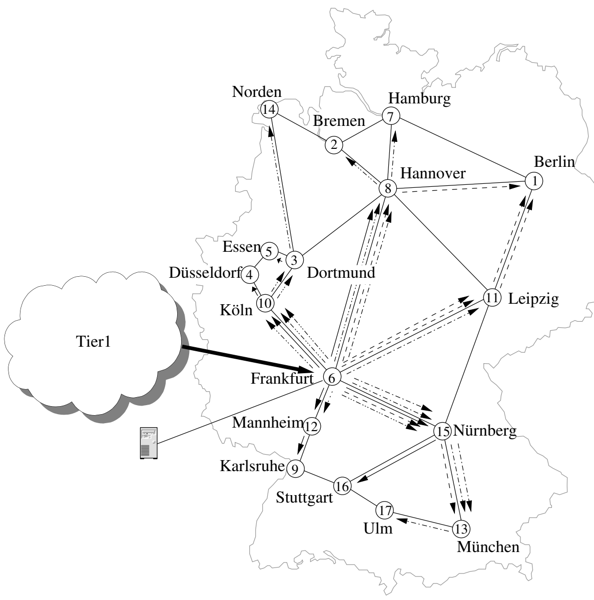

Figure 6: Traffic pattern by user generated content

In an announcement on July 2006 [17], 100 million clips were viewed daily on YouTube, with an additional 65,000 new videos uploaded per 24 hours. This translates to a total outbound traffic of over $40\ \mathrm{Gbit/s}$. Such a demand in bandwidth, naturally leads to schemes which reduce the strain on the network. Deployment of Content Delivery Networks (CDNs) results in bandwidth savings, as will be analyzed. In the following the total traffic volume of three major video sharing websites (YouTube, MySpace Videos, and Yahoo! Video) is estimated. Comparing the respective values of the market share of visits of the aforementioned sites in the US and globally, it is observed that they are very similar. In specific, YouTube has 47.7% in the US and 45.46% globally; MySpace Videos has 24.81% in the US and 22.99% globally. Finally, Yahoo! Video Search has 6.85% in the US and 6.06% globally [16].

When examining traffic values, what is of great importance is not the market share, but the actual initiated streams. At this point, the assumption is made that the average streams per streamer of the US user is representative of the global user and that the average size per stream is constant for all sites. This assumption is based on the estimation that the US user is representative for a global user in the market share case, as shown previously. Knowing that the outbound traffic of YouTube is over 40 Gbit/s and using the values for average streams per US streamer, the total outbound traffic of the three major video sharing websites is estimated to over 180 Gbit/s [18].

The growth of YouTube is estimated to be about 20% per month, without saturation until 2010. This has tremendous bandwidth costs and alternative delivery schemes are examined. Video Sharing Sites in general are experiencing traffic growth, perhaps at a slightly slower rate than YouTube. Nevertheless, they have emerged into an important traffic generating application in the Internet, which demands the deployment of schemes that reduce the bandwidth strain and the related costs. CDNs can result in significant bandwidth savings.

CDN is a term that describes a system of networked computers across the Internet, which delivers content to end users by cooperating transparently. In order to optimize the delivery process according to a set objective, content is moved between CDN nodes. The optimization objective can be the minimization of the bandwidth costs given a required end-user performance. The content type delivered is usually large media content.

CDNs augment the end-to-end transport network by distributing on it a variety of intelligent applications, employing techniques designed to optimize content delivery. The resulting tightly integrated overlay uses web caching, server-load balancing, request routing, and content services [19]. The number of nodes and servers making up a CDN depends on the architecture, some reaching thousands of nodes with tens of thousands of servers. Requests for content are served by these nodes that can serve content quickly to the user. The speed of delivery can be measured by the number of hops or the number of network seconds from the user. Another factor in cost, are the nodes that are less expensive to serve from. It is often the case that the goals of high speed and low cost are not conflicting, as servers that are close to the end user tend to have lower serving costs, since they probably are located within the same network as the end user.

As a case study, the CDN of Akamai will be examined in this section. Akamai now controls well over half the content distribution market. It is one of the world's largest on-demand distributed computing platform, with more than 20,000 servers in nearly 1,000 networks in 71 countries [20].

The daily web traffic carried by Akamai is greater than a Tier-1 ISP - at times reaching 200 Gbit/s. For example, during a sport event in April 2006, Akamai's traffic rate peaked above 200 Gbit/s serving 400,000 simultaneous video streams in a single day [21]. The rising need in delivery of rich media content (e.g. video) provides necessary conditions for growth in the sector of CDNs, to meet the needs of the rising market. For the examined reference model, the estimated traffic value is 5,555 Gbit/s.

3.5 Virtual Private Networks

A virtual private network (VPN) is a private communications network that runs over a publicly accessible network. Its main application is to serve communication purposes of companies and organizations that wish confidentiality and security, enabling them to extend their network service to branch offices and strategic partners. The global presence of the Internet has been one of the driving forces of the growth of VPNs.

VPNs are advantageous compared with dedicated private lines, since they provide a more cost-efficient and scalable solution. VPN traffic can be carried over a service provider's private network with a defined Service Level Agreement (SLA) or over a public network. The shared service provider backbone network is known as the VPN backbone and is used to transport traffic for multiple VPNs, as well as possibly non-VPN traffic. VPNs with technologies such as Frame Relay and Asynchronous Transfer Mode (ATM) virtual circuits (VC) have been available for a long time, but over the past few years IP and IP/Multiprotocol Label Switching (MPLS)-based VPNs have become more and more popular [22]. VPNs can be classified into site-to-site and remote access VPNs, independent on whether they are provider or customer provisioned. Site-to-site VPNs allow connectivity between an organizations' geographically dispersed sites. Remote access VPNs allow mobile and home-based users to access resources of organisations remotely.

The elimination of the need for expensive long-distance leased lines along with the cost reduction for backbone equipment and operations is leading to a migration of traditional services to VPN services. According to industry research, site-to-site connectivity costs are typically reduced by average 30% over domestic leased line networks. Cost reduction for client to site dial access is even greater, in the 60% to 80% range [23]. As a result a significant increase in the VPN traffic is expected in the coming years.

Previous studies [4] have shown that the expected Compound Annual Growth Rate (CAGR) for business traffic is in the range of 30% to 45%. Taking into account the growing market share of VPNs, the expected growth rate is even higher. The above is in accordance with conducted studies concerning the forecasted revenues of VPN services [24] as well as of Ethernet services [25]. For the examined reference model, the estimated value for the VPN traffic is 1,024 Gbit/s.

4 Results

In this section the traffic demand between nodes for each of the previously analyzed services is presented. A service oriented traffic model was constructed that follows the individual characteristics of every service as previously presented. This model served as the basis for the graphs depicted in **Figure 7** and **Figure 8**.

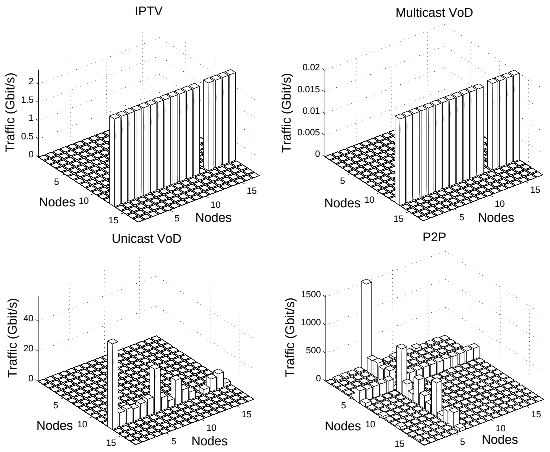

Figure 7: Traffic volume per service

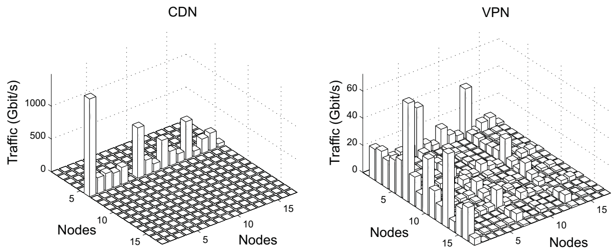

Figure 8: Traffic volume per service

In **Figure 9** the worst case unicast VoD traffic is also included, although there is currently, to the best of our knowledge, no such offering in the German market. However, when constructing the graph for the aggregated traffic, its produced traffic was omitted.

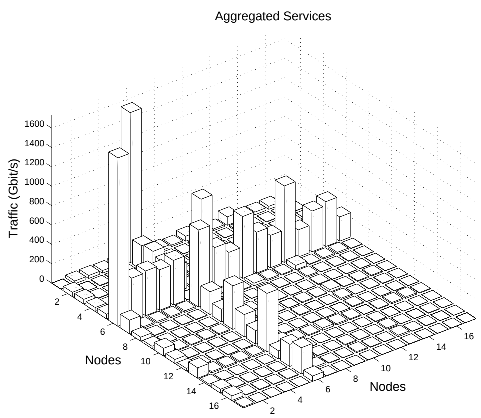

Figure 9: Aggregated traffic volume for all services

By comparing the different graphs, it is shown that the traffic demand between nodes is quasi-symmetric. The mean relative deviation of a demand pair (m,n) and (n,m) from their average is only 14%. The services that are dominating in terms of traffic are P2P and user-generated content, which is included in the CDN graph. In the case of multicast distribution schemes, the required bandwidth to carry IPTV and VoD traffic is almost negligible. Unicast VoD traffic with a low hit-ratio at the VHOs leads to enormous traffic demands.

5 Conclusion

This paper has discussed the traffic generated by existing and emerging services, which are expected to considerably contribute to the overall traffic in backbones. The set of considered services are IP Television (IPTV), Video on Demand (VoD) with content either fully distributed or partially distributed, Peer-to-Peer (P2P) applications, User Generated Content, Content Delivery Network (CDN), and Virtual Private Network (VPN).

If the partially-distributed variant of VoD is realized by unicast flows in the backbone, the hit-ratio for finding videos in the distributed content becomes a sensitive parameter. A low hit-ratio causes excessive backbone traffic to distribute videos from central server sites.

P2P, user generated content, and CDN dominate the overall traffic. The traffic is mainly star-oriented, i.e., the predominant traffic flows from and to central sites (such as server sites and peering points). In our modeling, traffic from IPTV, VoD (except for the case above), and VPN, even without considering oversubscribing, can be neglected in the backbone.

Future work includes studies on American and Asian networks, investigation of multiple central sites, and extending the models for VoD and VPN.

Acknowledgement

This work was supported in part by the German ministry for education and research (BMBF) under Grant 01BP551 (EIBONE). Responsibility for the content lies with the authors.

References

[1] Nathaniel Leibowitz, Aviv Bergman, Roy Ben-Shaul, and Aviv Shavit. Are file swapping networks cacheable? Characterizing P2P traffic. In *Proc. of the 7th Int. WWW Caching Workshop*, August 2002.

[2] Anurag Dwivedi and Richard E. Wagner. Traffic model for USA long-distance optical network. In *Optical Fiber Communication Conference*, March 2000.

[3] Chuck Fraleigh, Sue Moon, Bryan Lyles, Chase Cotton, Mujahid Khan, Deb Moll, Rob Rockell, Ted Seely, and Christophe Diot. Packet-level traffic measurements from the Sprint IP Backbone. *IEEE Network*, 17:6–16, 2003.

[4] Ralf Hülsermann, Stefan Bodamer, Marc Barry, Andreas Betker, Christoph Gauger, Monika Jäger, Martin Köhn, and Jan Späth. A set of typical transport network scenarios for network

modelling. In *5. ITG-Fachtagung Photonische Netze*, pages 65–72, May 2004.

[5] Meeyoung Cha, Gagan Choudhury, Jennifer Yates, Aman Shaikh, and Sue Moon. Case study: Resilient backbone design for IPTV services. In *Workshop on IPTV Services over World Wide Web*, May 2006.

[6] High quality and resilient IPTV multicast architecture - an overview of RESIP multicast architecture design guidelines. Technical report, Siemens Communications and Juniper Networks, Inc., 2006.

[7] Gartner. Press release. Online: `http://www.gartner.com/press_releases/asset_148795_11.html`, April 2006.

[8] D. Erman, D. Ilie, and A. Popescu. Bittorrent session characteristics and models. In *Proceedings of the 3rd International Working Conference on Performance Modelling and Evaluation of Heterogeneous Networks*, July 2005.

[9] LightReading. Broadband traffic cop needed. Online: `http://www.lightreading.com/document.asp?doc_id=103287&print=true`, September 2006.

[10] Andrew Parker. P2P media summit - Addressing the cost and performance challenges of digital media content delivery. Online: `dcia.info/P2PMSLA/CacheLogic.ppt`, 2006.

[11] ipoque GmbH. Press release. Online: `http://www.ipoque.com/dt/pressemitteilung_ipoque_231006.html`, October 2006.

[12] CacheLogic. Peer-to-Peer in 2005. Online: `http://www.cachelogic.com/home/pages/research/p2p2005.php`, March 2006.

[13] Krishna P. Gummadi, Richard J. Dunn, Stefan Saroiu, Steven D. Gribble, Henry M. Levy, and John Zahorjan. Measurement, modeling, and analysis of a peer-to-peer file-sharing workload. In *Proceedings of the nineteenth ACM symposium on Operating systems principles*, 2003.

[14] T. Karagiannis, P. Rodriguez, and K. Papagiannaki. Should internet service providers fear peer-assisted content distribution? In *In Proceedings of Internet Measurement Conference*, October 2005.

[15] James Enck. DAIWA institute of research LTD.: Eurotelco snapshot - Thinking about the data on data. Online: `http://193.194.136.7/news/archive/`, April 2006.

[16] LightReading. Top ten video sharing websites. Online: `http://www.lightreading.com/document.asp?doc_id=100934&print=true`, August 2006.

[17] YouTube. Fact sheet - Overview and features. Online: `http://www.youtube.com/t/fact_sheet`, 2007.

[18] Comscore Networks. Press release. Online: `http://www.comscore.com/press/release.asp?press=1015`, September 2006.

[19] Markus Hofmann and Leland R. Beaumont. *Content Networking: Architecture, Protocols, and Practice*. Morgan Kaufmann Publishers, 2005.

[20] Akamai. Online: `http://www.akamai.com/`.

[21] Rich Miller. Akamai traffic tops 200 Gigabits Per Second. Online: `http://www.datacenterknowledge.com/archives/2006/Mar/21/akamai_traffic_tops_200_gigabits_per_second.html`, March 2006.

[22] Cisco. What is a virtual private network. Online: `http://www.ciscopress.com/content/images/1587051796/samplechapter/1587051796content.pdf`.

[23] VPN-Info. Virtual private network guide. Online: `http://www.vpn-info.com/vpn_tunnel.htm`, 2007.

[24] Infonetics Research. Forecast for VPNs from infonetics VPN and security services report. Online: `http://www.infonetics.com/resources/purple.shtml?ms06.vs.nr.shtml`, July 2006.

[25] Infonetics Research. Forecast for ethernet services from infonetics ethernet services report. Online: `http://www.infonetics.com/resources/purple.shtml?ms06.eth.1q.nr.shtml`, April 2006.